

CLAIMS

1. A method of storing a portion of a packet in a host computer memory, comprising:
 - receiving a first packet at a communication interface;
 - receiving a second packet at said communication interface;
 - storing a header portion of said first packet in a hybrid storage area of a host computer;
 - if said first packet includes a data portion, storing said data portion in a data storage area of said host computer; and
 - if said second packet is smaller than a predetermined size, storing said second packet in said hybrid storage area.
2. The method of claim 1, wherein said hybrid storage area is substantially equal in size to one memory page of said host computer.
3. The method of claim 1, wherein said predetermined size is approximately 256 bytes.
4. The method of claim 1, further comprising:
 - receiving a code associated with said first packet, said code indicating that a data portion of said first packet is re-assembleable;
 - wherein said data storage area is a re-assembly data storage area.
5. The method of claim 1, further comprising:
 - receiving a code associated with said second packet, said code indicating that a data portion of said second packet is not re-assembleable; wherein said data storage area is a non-re-assembly data storage area.
6. The method of claim 1, further comprising receiving a code with said first packet, said code indicating that said first packet does not contain a data portion.

7. The method of claim 1, further comprising receiving a code with said second packet, said code indicating that said second packet is smaller than said predetermined size.

8. The method of claim 1, further comprising padding said hybrid storage area to align said header portion with a predetermined index in said hybrid storage area.

9. The method of claim 1, wherein said first packet belongs to a first communication flow and said second packet belongs to a second communication flow.

10. A method of transferring multiple packets in a communication flow into a host computer, comprising:

receiving a first packet at a network interface for transfer to a host computer;

identifying a communication flow comprising said first packet;

receiving a second packet at said network interface, wherein said communication flow further

comprises said second packet;

associating a code with said first packet to indicate whether a data portion of said first packet is re-assembleable with a data portion of another packet in said communication flow;

receiving a set of descriptors from the host computer, wherein an aggregate size of said set of descriptors approximates a cache line size of the host computer;

storing a header portion of said first packet in a header storage area;

storing said data portion of said first packet in a re-assembly storage area; and

storing a data portion of said second packet in said re-assembly storage area.

11. The method of claim 10, wherein said re-assembly storage area is substantially equal in size to one memory page of said host computer.

12. The method of claim 10, wherein said re-assembly storage area is used to store only data portions of packets in said communication flow.

13. The method of claim 10, further comprising retrieving an identifier of a first storage area in a host computer memory.

14. The method of claim 13, further comprising:

placing said first storage area identifier in a data structure configured to hold storage area identifiers;

wherein said first storage area identifier is identifiable by an index in said data structure.

15. The method of claim 14, further comprising using said index to identify said first storage area for storing a portion of said first packet.

16. The method of claim 14, wherein said first storage area comprises one of said header storage area and said re-assembly storage area, the method further comprising:

configuring a descriptor to store said index of said first storage area identifier to inform said host computer of the use of said first storage area to store one of said header portion and said data portion.

17. A method of network communication, the method comprising:

providing a computer having a processor and a memory, the memory including first, second and third storage areas that are accessible by a communication interface;

receiving, by the communication interface, a first packet, and storing the first packet in the first storage area;

receiving, by the communication interface, a second packet, storing a header of the second packet in the first storage area, and storing data of the second packet in the second storage area; and

receiving, by the communication interface, a third packet, and storing data of the third packet in the third storage area, the third storage area containing data from a plurality of packets and corresponding to an application running on the computer, the third storage area containing no headers.

18. The method of claim 9, further comprising parsing a header portion of said first packet.
19. The method of claim 18, wherein said parsing comprises:
 - retrieving an identifier of a source of said first packet; and
 - retrieving an identifier of a destination of said first packet.
20. The method of claim 19, wherein said identifying a communication flow comprises assembling said source identifier and said destination identifier to form a flow key.
21. The method of claim 20, said identifying further comprising obtaining an index of said flow key in a database of flow keys.
22. The method of claim 9, wherein said identifying comprises:
 - receiving a virtual connection identifier, wherein a communication flow comprising said first packet can be identified by said virtual connection identifier.
23. The method of claim 9, wherein said storing a header portion comprises:
 - retrieving an index of a header storage area of a host computer;
 - retrieving an identifier of a location in said header storage area; and
 - storing a header portion of said first packet at said location in said header storage area.
24. The method of claim 23, wherein said header storage area index comprises an index of said header storage area within a collection of storage area identifiers.
25. The method of claim 24, wherein said identifier of a location in said header storage area comprises an address within said header storage area.
26. The method of claim 9, wherein said storing said data portion comprises:
 - retrieving an identifier of a re-assembly storage area of a host computer;

retrieving an identifier of a location in said re-assembly storage area; and
storing said data portion of said first packet at said location in said re-assembly storage area.

27. The method of claim 26, wherein said retrieving an identifier of a re-assembly storage area comprises reading an entry in a re-assembly storage table, said entry corresponding to said communication flow.

28. The method of claim 18, wherein said data portion is of unknown size prior to said parsing.

29. A network interface for transferring a packet from a network to a host computer, comprising:

- a parser configured to examine one or more packets received from a network;
- a flow manager configured to manage a first communication flow comprising said one or more packets;
- a header storage area configured to store header portions of said one or more packets;
- a first re-assembly storage area configured to store data portions of said one or more packets;
- a second re-assembly storage area; and
- a transfer module configured to:
 - transfer a data portion of a first packet of a second communication flow into said second re-assembly storage area if said first packet is larger than a predetermined size; and
 - otherwise, transfer said first packet into said header storage area.

30. A computer readable storage medium storing instructions that, when executed by a computer, cause the computer to perform a method of transferring a packet from a network to a host computer, the method comprising:

- receiving a first packet at a communication interface;

receiving a second packet at said communication interface;
storing a header portion of said first packet in a hybrid storage area of a host computer;
if said first packet includes a data portion, storing said data portion in a data storage area of said host computer; and
if said second packet is smaller than a predetermined size, storing said second packet in said hybrid storage area.

31. A method of transferring a packet from a communication link to a host computer, comprising:

parsing a first packet received from a communication link to determine if the first packet conforms to a predetermined set of communication protocols;
generating, from identifiers of a source and a destination of the first packet extracted from the first packet, a flow key to identify a communication flow comprising the first packet;

updating a flow database to track a status of the communication flow;
associating a first operation code with the first packet to indicate whether the first packet is re-assembleable with another packet in the communication flow;
maintaining a plurality of storage areas for transferring packets to a host computer, including:

a first storage area configured to store data portions of packets in the communication flow;

a second storage area configured to store a packet having an associated operation code indicating that the packet is not re-assembleable; and

a third storage area configured to store packets less than a predetermined size and headers of packets having data portions stored in said first storage area; and

storing the first packet in one or more of said plurality of storage areas for transfer to the host computer.

32. The method of claim 31, wherein said storing comprises:
selecting one or more storage areas of said plurality of storage areas on the basis of said operation code.
33. The method of claim 31, wherein each of said plurality of storage areas is substantially equal in size to one memory page of the host computer.
34. The method of claim 17, further comprising associating the third storage area with a transmission control protocol (TCP) connection.
35. The network interface of claim 29, further comprising:
a code generator configured to generate a code for the first packet to indicate whether the first packet is re-assembleable with a second packet in the second communication flow.
36. The network interface of claim 29, further comprising:
a code generator configured to generate a code for the first packet to indicate whether the first packet is larger than the predetermined size.
37. A communication interface configured to store packets for transfer to a host computer, comprising:
a parser configured to determine whether a packet includes a data portion;
a re-assembly storage area configured to store data portions of a plurality of packets from a single communication flow; and
a hybrid storage area configured to store:
header portions of the plurality of packets; and
one or more packets smaller than a first predetermined size.
38. The method of claim 17, further comprising sending a command from the computer to the communication interface, the command including an address of the third storage area.

39. The communication interface of claim 37, further comprising:
a code generator configured to generate a code for each packet received at the communication interface;
wherein said code is configured to indicate a type of storage area in which said packet may be stored.
40. The communication interface of claim 37, wherein said parser is further configured to determine a size of each packet received at the communication interface.